



Attorney Docket No.: NCR 10031
Express Mail Label No.: _____ US
Date of Deposit with Express Mail: _____, 2004

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Paul L. Sinclair, et al.	§	Group Art Unit: 2171
	§	
Serial No.: 10/017,406	§	
	§	
Filed: December 7, 2001	§	Examiner: C. Nguyen
	§	
Title: Reorganization of Database Records Using Partition Merging	§	
	§	Attorney Docket No: NCR 10031

DRAFT

Honorable Commissioner for Patents
Washington, D.C. 20231
Attention: Board of Patent Appeals and Interferences

Dear Sir:

APPELLANT'S BRIEF (37 C.F.R. § 1.192)

This brief is submitted in support of appellants' notice of appeal from the decision dated April 23, 2004 of the Examiner.

This brief is transmitted in triplicate per 37 C.F.R. § 1.192.

I. IDENTIFICATION OF THE REAL PARTY OF INTEREST

The real party in interest is:

NCR Corporation
1700 S. Patterson Blvd
Dayton, Ohio 45479

by virtue of an assignment by the inventors as duly recorded in the Assignment Branch of the U.S. Patent and Trademark Office.

II. IDENTIFICATION OF RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences, to Applicants' knowledge.

III. STATUS OF ALL THE CLAIMS, PENDING OR CANCELLED, AND IDENTIFYING THE CLAIMS APPEALED

The application as originally filed contained 53 claims. Claims 1-53 are pending. Claims 1-12, 14-27, 29-43, 45-51, and 53 are appealed. Claims 13, 28, 44 and 52 have not been rejected and have been identified as allowable subject matter by the Examiner.

IV. STATUS OF ANY AMENDMENT FILED SUBSEQUENT TO FINAL REJECTION

No amendment has been filed subsequent to final rejection.

V. SUMMARY OF THE INVENTION

Applicant's invention is a method of reorganizing rows of a database table that are initially stored in partitions. Within each partition that contains at least one row (each populated partition) the rows are stored in order of a first value associated with each row. For example, in a table of sale transactions, the rows could be partitioned by the month in which the sale occurred and ordered within the month-based partitions by the result of applying a hash function to a sales transaction number, that result is called the hash value. A subset of the populated partitions are then selected and a file context is created for each partition in the subset. Each file context stores location data for a row in its partition and the first value for that row. To continue the example, the first three months in which sales occurred could be selected and three file contexts would be created. Each of those three file contexts could store the location of the row for that month that has the lowest hash value as well as the hash value itself. The partitions in the subset are then merged into a single partition with the rows in order of the first value. In the example, the rows for each of the sales made in the first three months would be ordered as a group according to the hash value. Those rows would no longer be ordered by the month in which the sale occurred. The steps of selecting a subset of populated partitions, creating file contexts, and merging the subset into one partition is repeated until all the original populated partitions have been merged into partitions that each correspond to a subset of the original populated partitions. In the

example, all the rows in the table of sales transactions could be stored in order of the hash value in three month increments.

Reorganizing the table rows so that they are no longer in partitions can be important to improving the performance of queries that do not include selection of rows according to the values by which the rows were partitioned, but do include a join of the table to another table based on the first value.

VI. CONCISE STATEMENT OF THE ISSUES PRESENTED FOR REVIEW

The examiner rejected claims 1-7, 12, 14, 15, 17-22, 27, 29, 30, 32-38, 43, 45, 46, and 48-51 under 35 USC 102(e) as being anticipated by Agarwal et al. (United States Patent Number 6,223,182).

With regard to independent claims 1, 17, 32, and 48, the examiner argued that Agarwal discloses creating a file context for each partition and that the disclosed file context stores location data and a first value for a row in that partition. In the Claim Rejections section, the examiner identified Agarwal's intermediate table as the source of this disclosure. In the Response to Arguments section, the examiner identified Agarwal's partition table as the source of this disclosure. The applicant disagrees that Agarwal makes such a disclosure with regard to either table.

With regard to dependent claims 3, 19, and 34, the examiner argued that Agarwal discloses using the location of a block containing the row as the location data for a row stored in a file context. The applicant disagrees that Agarwal makes such a disclosure.

With regard to dependent claims 6, 21 and 37, the examiner argued that Agarwal discloses storing the rows of the merged partitions separately from the rows of the populated partitions. The applicant disagrees that Agarwal makes such a disclosure.

With regard to dependent claims 7, 22, 38, and 50, the examiner argued that Agarwal discloses determining whether rows from a partitioned primary index table are being spooled and whether a subsequent operation requires the spooled rows to be ordered in accordance with the first value, and performing the partition merger operation only if the two determinations are true. The applicant disagrees that Agarwal makes such a disclosure.

The examiner rejected claims 8-11, 23-26, 31, 39-42, 47 and 53 under 35 USC 103(a) as

being unpatentable over Agarwal in view of Goetze Graefe, "Query Evolution techniques for large databases", ACM computing surveys, Vol. 25, No. 2, June 1993. The examiner argued that Agarwal discloses the limitations of claims rejected under 102(e) from which the claims rejected under 103(a) depend. The applicant disagrees that Agarwal makes such a disclosure.

VII. GROUPING OF CLAIMS

Claims 1-2, 4-5, 8-13, 14-18, 20, 23-29, 30-33, 35-36, 39-45, 46-49, and 51-53 stand or fall together.

Claims 3, 19, and 34 stand or fall together.

Claims 6, 21 and 37 stand or fall together.

Claims 7, 22, 38, and 50 stand or fall together.

VIII. ARGUMENTS OF THE APPELLANT, WITH EACH ISSUE IN SEPARATE HEADINGS, WITH RESPECT TO EACH ISSUE PRESENTED FOR REVIEW

- A. Agarwal does not anticipate or render obvious claims 1-12, 14-27, 29-43, 45-51, and 53 because Agarwal does not disclose creating a file context for each partition in a subset of the populated partitions that stores location data for a row and a first value for that row, as required by claims 1-12, 14-27, 29-43, 45-51, and 53.**

In the Final Office Action in this matter, the examiner rejected claims 1, 17, 32, and 48 under 35 USC 102(e) as being anticipated by Agarwal et al. (United States Patent Number 6,223,182). Claims 1, 17, 32, and 48 each require, in pertinent part, "creating a file context for each partition of a subset of the populated partitions, each file context storing at least location data for a row in the partition and the first value associated with the row." Claims 2-12, 14-16, 18-27, 29-31, 33-43, 45-47, 49-51, and 53 each depend from one of claims 1, 17, 32, and 48.

In the rejection of claims 1-12, 14-27, 29-43, 45-51, and 53, the examiner argued that Agarwal discloses "creating a file context (intermediate table) for each partition of a subset of the populated partitions (col. 6, lines 4-11, Agarwal), each file context storing at least location data for a row in the partition and the first value associated with the row (col. 6, lines 12-19, Agarwal)." See page 3 of the Final Office Action.

The intermediate table discussed in the portions of Agarwal cited by the examiner, however, only includes row-specific information, not partition-specific information: "an

intermediate table containing BH codes and rowids of the original table or view is created within the database.” (col. 6, lines 6-8). Figure 4 of Agarwal depicts an example intermediate table that contains a row for each rowid. None of the information in Figure 4 corresponds to a particular partition. In fact, the intermediate table cannot include partition-specific information because the partitions are not generated until after the intermediate table is created. “The intermediate table allows for database operations to be performed to determine how to partition the table without causing the data in the table to go offline.” (col. 6, lines 8-11). Figure 2 of Agarwal depicts this sequence by showing the creation of the intermediate table at step 202 and the generation of partitions later at step 216. As discussed in Agarwal (col. 8, lines 20-30), the intermediate table is merely a copy of information from the table prior to partitioning that allows the table to be accessed during the partitioning process. The intermediate table does not include the claimed file contexts for each partition in a subset of partitions.

In responding to applicants’ arguments, the examiner offered a new argument that partition table 1000 (illustrated in Figure 10 and discussed at col. 12, lines 30-36), rather than the intermediate table, disclosed the partition-specific file contexts required by the claims. The partition table 1000 does include partition-specific information, but not the partition-specific information required by the claims. The claimed file contexts must store “at least location data for a row in the partition and the first value associated with the row.” The partition table 1000 includes neither type of information. Instead, the partition table merely stores the partition identifier (P_CCODE), the current number of rows in the partition (P_COUNT), and the number of rows to be added to the partition (P_NEW). The partition table does not include the claimed file contexts for each partition in a subset of partitions. Anticipation can be established only when every element of the claim is disclosed by a single prior art reference. MPEP 2131; *RCA Corp. v. Applied Digital Data Systems, Inc.*, 730 F.2d 1440, 1444, 221 USPQ 385, 388 (Fed. Cir. 1984). The rejection of claims 1-7, 12, 14, 15, 17-22, 27, 29, 30, 32-38, 43, 45, 46, and 48-51 under 35 USC 102(e) should be withdrawn because Agarwal does not disclose creating the claimed file context for each partition of a subset of partitions. There is no prima facie case of obviousness where the asserted combination lacks at least one element. MPEP 2143; *In re Vaeck*, 947 F.2d 488, 493, 20 USPQ2d 1438, 1443 (Fed. Cir. 1991). The rejection of claims 8-11, 23-26, 31, 39-42, 47 and 53 under 103(a) should be withdrawn because the asserted

combinations do not disclose creating the claimed file context for each partition of a subset of partitions.

B. Agarwal does not anticipate claims 3, 19, and 34 because Agarwal does not disclose using the location of a block containing the row as the location data for a row stored in a file context, as required by claims 3, 19, and 34.

The examiner rejected claims 3, 19, and 34 as being anticipated by Agarwal et al. (United States Patent Number 6,223,182). Claims 3, 19, and 34 each require, in pertinent part, "the location data for a row is the location of a block of rows that includes the row." As discussed above, the independent claims from which claims 3, 19, and 34 depend require that a file context be created for each of a subset of populated partitions and that each file context include location data for a row in its partition.

In rejecting claims 3, 19, and 34, the examiner argued that Agarwal discloses "wherein the location data for a row in the location of a block of rows that includes the row (col. 6, lines 12-19, Agarwal)." See page 4 of the Final Office Action. The portion of Agarwal discussed by the examiner discusses an index-organized table within which the rows of the table are sorted and stored. This is a different table from either the intermediate table or the partition table that the examiner relied upon as disclosing the file contexts. Even assuming that the index-organized table were relevant to the information contained in the intermediate table and the partition table, there is no disclosure that the index-organized table includes the location of a block of rows containing a row from a partition. Agarwal's only discussion of blocks does not teach that the location of blocks is stored, rather "the rows are sorted and stored based on the BH codes so that rows that are close together (as defined by the BH codes) are likely to be located in the same disk block." (col. 6, lines 15-18). While rows are likely to be located in the same disk block, there is no disclosure of storing that information with a file context. None of the tables identified by the examiner as potentially disclosing the file context (Figures 4 and 10) include block location information. Anticipation can be established only when every element of the claim is disclosed by a single prior art reference. MPEP 2131; *RCA Corp. v. Applied Digital Data Systems, Inc.*, 730 F.2d 1440, 1444, 221 USPQ 385, 388 (Fed. Cir. 1984). The rejection of claims 3, 19, and 34 under 35 USC 102(e) should be withdrawn because even if Agarwal disclosed creating the claimed file context for each partition of a subset of partitions, it does not disclose including block location information in each file context.

C. Agarwal does not anticipate claims 6, 21 and 37 because Agarwal does not disclose using the location of a block containing the row as the location data for a row stored in a file context, as required by claims 6, 21 and 37.

The examiner rejected claims 6, 21 and 37 as being anticipated by Agarwal et al. (United States Patent Number 6,223,182). Claims 6, 21 and 37 each require, in pertinent part, "the rows of the first-merge partitions are stored separately from the rows of the populated partitions of the partitioned database table."

In rejecting claims 6, 21 and 37, the examiner argued that Agarwal discloses "wherein the rows of the first-merge partitions are stored separately from the rows of the populated partitions of the partitioned database table (col. 11, lines 64 to col. 12, lines 2, Agarwal)." See page 5 of the Final Office Action. The portion of Agarwal discussed by the examiner proposes an index-organized table in which ranges specified by the branch nodes correspond to the proposed partitions identified in steps 200-214. The index-organized table is an alternate storage configuration: "In one embodiment, a data container is organized, not by partitioning the data in the sense of storing each subset of the data container in separate datafiles as in other embodiments, but by storing the data container as an index-organized table" (col. 11, lines 64-67) (emphasis added), it does not disclose separately storing table rows in two different sets of partitions. Agarwal teaches use of only one of two organizations, not separate storage in two different organizations. Anticipation can be established only when every element of the claim is disclosed by a single prior art reference. MPEP 2131; *RCA Corp. v. Applied Digital Data Systems, Inc.*, 730 F.2d 1440, 1444, 221 USPQ 385, 388 (Fed. Cir. 1984). The rejection of claims 6, 21 and 37 under 35 USC 102(e) should be withdrawn because even if Agarwal disclosed creating the claimed file context for each partition of a subset of partitions, it does not disclose separate row storage of merge partitions and original partitions.

D. Agarwal does not anticipate claims 7, 22, 38, and 50 because Agarwal does not disclose determining whether rows from a partitioned primary index table are being spooled and whether a subsequent operation requires the spooled rows to be ordered in accordance with the first value, and performing the partition merger operation only if the two determinations are true, as required by claims 7, 22, 38, and 50.

The examiner rejected claims 7, 22, 38, and 50 as being anticipated by Agarwal et al. (United States Patent Number 6,223,182). Claims 7, 22, 38, and 50 each require, in pertinent part, "determining whether rows from a partitioned primary index table are being spooled,"

“determining whether a subsequent operation requires the spooled rows to be ordered in accordance with the first value associated with each row,” and selecting and merging partition subsets “only if both determinations [] are true.”

In rejecting claims 7, 22, 38, and 50, the examiner argued that Agarwal discloses:

“a'. determining whether rows from a partitioned primary index table are being spooled (col. 11, lines 1-38, Agarwal);

a". determining whether a subsequent operation requires the spooled rows to be ordered in accordance with the first value associated with each row (col. 11, lines 18-38, Agarwal); and

a"". performing steps b through d only if both determinations, a' and a", are true (col. 13, lines 21-30, Agarwal).”

See page 5 of the Final Office Action. The portion of Agarwal discussed by the examiner does not discuss spooling of rows or any subsequent operations. Instead, it discloses that “partitions can be stored in a variety of formats” and discusses, in particular, a tree data structure with reference to Figure 7. While claims 7, 22, 38, and 50 require an “only if” determination of subsequent operations to determine whether partitions should be merged, the section of Agarwal referenced by the examiner states that “If the partitions can be coalesced according to the MAX parameter, then they are collapsed.” (col. 13, lines 26-27). Agarwal teaches away from merging partitions “only if” particular subsequent operations will occur. Anticipation can be established only when every element of the claim is disclosed by a single prior art reference. MPEP 2131; *RCA Corp. v. Applied Digital Data Systems, Inc.*, 730 F.2d 1440, 1444, 221 USPQ 385, 388 (Fed. Cir. 1984). The rejection of claims 67, 22, 38, and 50 under 35 USC 102(e) should be withdrawn because even if Agarwal disclosed creating the claimed file context for each partition of a subset of partitions, it does not disclose predicating the merger operation on the nature of subsequent operations.

Summary

The rejection of the pending claims should be reversed because none of the cited references, alone or in combination, disclose the elements required by the claims.

Respectfully submitted,

By: _____
Howard L. Speight
Registration No. 37,733
ATTORNEY FOR APPLICANT

Baker Botts L.L.P.
One Shell Plaza
910 Louisiana
Houston, Texas 77002-4995
Telephone: (713) 229-2057
Fax: (713) 229-2757

Date: _____, 2002